

Express Mail Label No EL 550 340 297 US

PATENT APPLICATION
Docket No. 13768.172

UNITED STATES PATENT APPLICATION

of

Jon B. Avner

and

Soner F. Terek

for

**METHODS AND SYSTEMS FOR
PERFORMING RELIABLE ASYNCHRONOUS NOTIFICATION OF
HIGH-LEVEL DOCUMENT OPERATIONS**

1

BACKGROUND OF THE INVENTION

2

1. The Field of the Invention

3

The present invention relates to the field of database management systems. In particular, the present invention relates to methods and systems for performing reliable asynchronous notification of high-level document operations in a database.

6

7

2. The Prior State of the Art

8

9

With the development of computer networks and the Internet, anyone who has access to an Internet capable computer may access information from all over the world. The present era has even been termed the "information age" due to the widespread abundance of information. Although this abundance of information is useful, individuals may easily be overwhelmed with information to the point where it is difficult to filter out relevant information from irrelevant information.

14

15

16

17

Database management systems are systems that manage collections of data (i.e., databases) so as to allow individuals, devices and applications to easily access, manage, and update information. Thus, database management systems provide a significant contribution to the information age by allowing for the efficient use of information.

18

19

20

21

22

23

24

Conventional database management systems such as the database management system 100 shown in Figure 1 include an underlying database 110 that contains organized data in the form of a number of tables such as table "A", table "B", table "C" and table "D". Each table contains entries that associate documents with properties of the documents. For example, one table may contain a count of the number of files within given folders, another may list the memory size of given files, yet another may list other properties associated with given files, or folders. For each document, there is typically

1 more than one table that may be associated with the document by, for example, describing
2 properties of the document.

3 A document (e.g., document 1, 2, 3 and 4 in Figure 1) is an identifiable entity from
4 the viewpoint of applications that use the database management system. For example, a
5 folder or an item such as an electronic mail message within the folder may be considered to
6 be a document since they are identified as entities from the viewpoint of the application
7 that uses the database management system. For example, in the context of electronic
8 messaging, an application may present folders such as "in-boxes" and "out-boxes" to a
9 user with corresponding electronic mail messages in the folders. These folders and
10 electronic mail messages are "documents" from the viewpoint of the application.
11 Similarly, files and directories within a file system may also be documents from the
12 viewpoint of the application that uses the database management system. The tables within
13 the underlying database are not considered to be documents since they are identified
14 internal to the database management system, and not at the higher level of the application
15 that uses the database management system.

16 The database application 120 generates high-level document commands (e.g., high-
17 level command 121) that relate to operations to be performed on a document. Examples of
18 such document commands might include operations such as, for example, move folder,
19 move message, delete message, copy folder, copy file, and so forth.

20 Each of these high-level document commands is received by the database engine
21 130 which implements the high-level document commands by executing a number of
22 object commands that result in objects such as tables being updated. Typically, the
23 database engine would use the disk access module 140 (e.g., a disk driver) of an operating
24 system to produce the physical control signals necessary to read and write the appropriate

1 sectors in the disk, each object comprising one or more possibly discontiguous sectors on
2 the disk.

3 It is often desirable for one client application to receive notification when a high-
4 level document command meeting certain criteria has been implemented by another client
5 application. For example, if a client application is viewing all electronic messages within a
6 folder, notification when another client applications adds a message to the folder will allow
7 the client application to refresh the display.

8 Figure 2 illustrates a flowchart of a conventional method 200 for performing
9 notifications that certain high-level document commands have been implemented in the
10 database. When receiving a given high-level document command, the database
11 management system first implements the corresponding high-level document command in
12 the database (act 201). Once implemented, the database management system determines
13 which applications have subscribed to be notified when the given high-level document
14 command is implemented (act 202). The database management system then dispatches a
15 notification that the high-level document command has been implemented (act 203).

16 The vast majority of the time, when a high-level document command having
17 certain criteria is issued, the subscribing client applications are properly notified of the
18 high-level document command and thus can respond appropriately. However, sometimes
19 computer systems become inoperative due to system failures, corruption, power loss, or the
20 like. In such cases, there is a possibility that the high-level document command is
21 implemented in the database, but that the subscribing client application never receives
22 notification of the implementation. For example, the client application may become
23 inoperative and thus become unable to receive and interpret a notification that a certain
24 high-level document command has been implemented. Also, the database management

1 system itself may become inoperative so that although the command has been
2 implemented, no notification was ever dispatched.

3 Thus, conventional system contain no guarantee that the proper notifications of a
4 high-level document command will always be dispatched or received if a high-level
5 document command is implemented. Therefore, what are desired are methods and systems
6 for improving accuracy and reliability in notifying client applications that certain high-
7 level document commands have been implemented in a database.

8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24

SUMMARY OF THE INVENTION

The principles of the present invention provide for the reliable asynchronous notification to client applications of the implementation of predefined high-level document commands that are implemented in a database managed by a database management system. These notifications are “asynchronous” since the high-level document commands are actually implemented in the database prior to sending the notification to client applications. Thus, the notified client applications do not affect the implementation of the high-level document commands in the database.

The database management system is issued a number of high-level document commands which are to be implemented in the underlying database. Such high-level document commands include high-level commands that relate to actions to be taken on documents such as folders, files, messages and other entities that are identified at the level of the application that uses the database management system. For example, in electronic messaging applications, one high-level document command may be to add a message to a folder.

16 Ultimately, the database engine implements the high-level document command by
17 altering one or more tables in the database. However, before implementing the high-level
18 command, an entry is set in a notification table stored in a persistent memory such as the
19 database. Alternatively, the entry in the notification table is set in the same transaction as
20 any other table changes that are made to implement the high-level command. In this
21 transaction, either all actions in the transaction are performed (all relevant tables are
22 altered) or no actions in the transaction are taken (none of the relevant tables are altered).
23 Thus, if the entry in the notification table is set, then the corresponding high-level
24 document command is guaranteed to have been implemented in the database.

Ideally, the client application that requested notification and the database management system will continue to be operational. In this ideal case, the notification that the given high-level document command has been implemented is dispatched to the client application. The client application then acknowledges to the database management system that the notification has been received. In response, the entry is deleted in the notification table in the database. The notification may be resent if the database management system does not receive such acknowledgement within an expected time period.

To make sure that the client management system properly performs all actions responsive to the notification, the entry in the notification table may remain set until the database management system receives acknowledgement from the client application that all necessary actions have been performed by the client application. Thus, if a client application fails while implementing any necessary actions responsive to the notification, the notification will be resent to allow the client application to re-perform any necessary actions once the client application becomes operational again.

Suppose now that a system failure occurs at the database management system after the high-level document command is implemented in the database, but before notification is dispatched. Upon restarting the database management system, the system will check the notification table to see what notifications entries are set. These entries correspond to notifications for which the database management system has not received acknowledgement from the client application. The database management system then sends corresponding notifications to the client application. In this manner, the client application receives notification even if a database management system failure occurs.

Now suppose that the notification is indeed dispatched, but that the client application has failed so as to be unable to receive the notification. The database

1 management system will not receive acknowledgement from the client application.
2 Therefore, the database application will resend the notification until the notification is
3 acknowledged upon restarting the client application. In this manner, the client application
4 receives notification even if a client application failure occurs.

5 Therefore, the principles of the present invention provide for a reliable way of
6 notifying client application when particular high-level document commands are
7 implemented in a database, the notifications surviving even failure of the database
8 management system or client application.

9 Additional features and advantages of the invention will be set forth in the
10 description which follows, and in part will be obvious from the description, or may be
11 learned by the practice of the invention. The features and advantages of the invention may
12 be realized and obtained by means of the instruments and combinations particularly
13 pointed out in the appended claims. These and other features of the present invention will
14 become more fully apparent from the following description and appended claims, or may
15 be learned by the practice of the invention as set forth hereinafter.

16
17
18
19
20
21
22
23
24

1 **BRIEF DESCRIPTION OF THE DRAWINGS**

2 In order to describe the manner in which the above-recited and other advantages
3 and features of the invention can be obtained, a more particular description of the invention
4 briefly described above will be rendered by reference to specific embodiments thereof
5 which are illustrated in the appended drawings. Understanding that these drawings depict
6 only typical embodiments of the invention and are not therefore to be considered to be
7 limiting of its scope, the invention will be described and explained with additional
8 specificity and detail through the use of the accompanying drawings in which:

9 Figure 1 illustrates a hierarchical view of a database management system in
10 accordance with the prior art;

11 Figure 2 illustrates a flowchart of a method of notifying client applications when a
12 particular high-level document command is implemented in accordance with the prior art;

13 Figure 3 illustrates an exemplary system that provides a suitable operating
14 environment for the present invention;

15 Figure 4 illustrates a hierarchical view of a database management system in
16 accordance with the present invention in which notifications survive system failures;

17 Figure 5 illustrates a flowchart of a method for guaranteeing notification despite
18 intervening system failures; and

19 Figure 6 illustrates a flowchart of a method for recovering from a system failure so
20 as to ensure notification transmission.

21
22
23
24

DETAILED DESCRIPTION OF THE INVENTION

The present invention extends to both methods and systems for reliably notifying client applications of the implementation of predefined high-level document commands in a database management system. These notifications are "asynchronous," meaning that the client applications do not affect the implementation of the high-level document command since the implementation takes place prior to the notification.

The notification occurs even if the database management system or the client application experiences a system failure. A notification table records implementation of high-level document commands. Then, notification is dispatched to the client applications that are to be notified. The client application then acknowledges at least receipt of the notification, and perhaps that they have completed performing actions in response to the notification. Upon receiving the acknowledgement, the notification table is updated to remove the record of the high-level document command.

If the database management system or the client application experiences a system failure that inhibits the notification process, the notification is resent after restarting the failed system. Upon restarting the database management system, for example, the database management system checks the notification table and resends any notifications for which there is a record (i.e., notifications which have not been acknowledged). Therefore, the present invention provides for the reliable notification of high-level document commands where the notifications survive even system failures.

The embodiments of the present invention may comprise a special purpose or general purpose computer including various computer hardware, as discussed in greater detail below. Embodiments within the scope of the present invention also include computer-readable media for carrying or having computer-executable instructions or data

1 structures stored thereon. Such computer-readable media can be any available media
2 which can be accessed by a general purpose or special purpose computer. By way of
3 example, and not limitation, such computer-readable media can comprise physical storage
4 media such as RAM, ROM, EEPROM, CD-ROM or other optical disk storage, magnetic
5 disk storage or other magnetic storage devices, or any other medium which can be used to
6 carry or store desired program code means in the form of computer-executable instructions
7 or data structures and which can be accessed by a general purpose or special purpose
8 computer.

9 When information is transferred or provided over a network or another
10 communications connection (either hardwired, wireless, or a combination of hardwired or
11 wireless) to a computer, the computer properly views the connection as a computer-
12 readable medium. Thus, any such a connection is properly termed a computer-readable
13 medium. Combinations of the above should also be included within the scope of
14 computer-readable media. Computer-executable instructions comprise, for example;
15 instructions and data which cause a general purpose computer, special purpose computer,
16 or special purpose processing device to perform a certain function or group of functions.

17 Figure 3 and the following discussion are intended to provide a brief, general
18 description of a suitable computing environment in which the invention may be
19 implemented. Although not required, the invention will be described in the general context
20 of computer-executable instructions, such as program modules, being executed by
21 computers in network environments. Generally, program modules include routines,
22 programs, objects, components, data structures, etc. that perform particular tasks or
23 implement particular abstract data types. Computer-executable instructions, associated
24 data structures, and program modules represent examples of the program code means for

1 executing steps of the methods disclosed herein. The particular sequence of such
2 executable instructions or associated data structures represent examples of corresponding
3 acts for implementing the functions described in such steps.

4 Those skilled in the art will appreciate that the invention may be practiced in
5 network computing environments with many types of computer system configurations,
6 including personal computers, hand-held devices, multi-processor systems,
7 microprocessor-based or programmable consumer electronics, network PCs,
8 minicomputers, mainframe computers, and the like. The invention may also be practiced
9 in distributed computing environments where tasks are performed by local and remote
10 processing devices that are linked (either by hardwired links, wireless links, or by a
11 combination of hardwired or wireless links) through a communications network. In a
12 distributed computing environment, program modules may be located in both local and
13 remote memory storage devices.

14 With reference to Figure 3, an exemplary system for implementing the invention
15 includes a general purpose computing device in the form of a conventional computer 320,
16 including a processing unit 321, a system memory 322, and a system bus 323 that couples
17 various system components including the system memory 322 to the processing unit 321.
18 The system bus 323 may be any of several types of bus structures including a memory bus
19 or memory controller, a peripheral bus, and a local bus using any of a variety of bus
20 architectures. The system memory includes read only memory (ROM) 324 and random
21 access memory (RAM) 325. A basic input/output system (BIOS) 326, containing the basic
22 routines that help transfer information between elements within the computer 320, such as
23 during start-up, may be stored in ROM 324.

24

1 The computer 320 may also include a magnetic hard disk drive 327 for reading
2 from and writing to a magnetic hard disk 339, a magnetic disk drive 328 for reading from
3 or writing to a removable magnetic disk 329, and an optical disk drive 330 for reading
4 from or writing to removable optical disk 331 such as a CD-ROM or other optical media.
5 The magnetic hard disk drive 327, magnetic disk drive 328, and optical disk drive 330 are
6 connected to the system bus 323 by a hard disk drive interface 332, a magnetic disk drive-
7 interface 333, and an optical drive interface 334, respectively. The drives and their
8 associated computer-readable media provide nonvolatile storage of computer-executable
9 instructions, data structures, program modules and other data for the computer 320.
10 Although the exemplary environment described herein employs a magnetic hard disk 339,
11 a removable magnetic disk 329 and a removable optical disk 331, other types of computer
12 readable media for storing data can be used, including magnetic cassettes, flash memory
13 cards, digital video disks, Bernoulli cartridges, RAMs, ROMs, and the like.

14 Program code means comprising one or more program modules may be stored on
15 the hard disk 339, magnetic disk 329, optical disk 331, ROM 324 or RAM 325, including
16 an operating system 335, one or more application programs 336, other program modules
17 337, and program data 338. A user may enter commands and information into the
18 computer 320 through keyboard 340, pointing device 342, or other input devices (not
19 shown), such as a microphone, joy stick, game pad, satellite dish, scanner, or the like.
20 These and other input devices are often connected to the processing unit 321 through a
21 serial port interface 346 coupled to system bus 323. Alternatively, the input devices may
22 be connected by other interfaces, such as a parallel port, a game port or a universal serial
23 bus (USB). A monitor 347 or another display device is also connected to system bus 323
24 via an interface, such as video adapter 348. In addition to the monitor, personal computers

1 typically include other peripheral output devices (not shown), such as speakers and
2 printers.

3 The computer 320 may operate in a networked environment using logical
4 connections to one or more remote computers, such as remote computers 349a and 349b.
5 Remote computers 349a and 349b may each be another personal computer, a server, a
6 router, a network PC, a peer device or other common network node, and typically include
7 many or all of the elements described above relative to the computer 320, although only
8 memory storage devices 350a and 350b and their associated application programs 336a and
9 336b have been illustrated in Figure 3. The logical connections depicted in Figure 3
10 include a local area network (LAN) 351 and a wide area network (WAN) 352 that are
11 presented here by way of example and not limitation. Such networking environments are
12 commonplace in office-wide or enterprise-wide computer networks, intranets and the
13 Internet.

14 When used in a LAN networking environment, the computer 320 is connected to
15 the local network 351 through a network interface or adapter 353. When used in a WAN
16 networking environment, the computer 320 may include a modem 354, a wireless link, or
17 other means for establishing communications over the wide area network 352, such as the
18 Internet. The modem 354, which may be internal or external, is connected to the system
19 bus 323 via the serial port interface 346. In a networked environment, program modules
20 depicted relative to the computer 320, or portions thereof, may be stored in the remote
21 memory storage device. It will be appreciated that the network connections shown are
22 exemplary and other means of establishing communications over wide area network 352
23 may be used.

24

Figure 4 illustrates a database management system 400 in accordance with the present invention. Although, not required, the database management system 400 may be implemented in the computing environment shown in Figure 4. In contrast to the prior art database management system shown in Figure 1, the database management system 400 includes a notification table 450. The operation of the database management system 400 of Figure 4 will be described with respect to the method 500 illustrated in Figure 5.

Figure 5 shows a flowchart of a method performed by the database management system 400 to perform reliable notification when a high-level document command is implemented by the database management system. The database application 420 issues a high-level document command 421 to the database engine 430. This high-level document command is for performing a function on a document that the database application recognizes to be a separately identifiable entity. For example, if the database management system supports electronic mail applications, the document may be an electronic mail message or a folder that contains the electronic mail message. Typical operations might be to move, delete, copy, update, or add a new document.

As the database engine 430 receives each high-level document command, the database engine implements the high-level document command in the database 410 (act 501). Typically, implementing a high-level document command in the database may involve changing entries in a number of tables such as tables "A", "B", "C" and "D" since each table could potentially have an entry for the corresponding document. The client application that uses the database management system 400 is typically not concerned with the details of which table needs to be changed. The database engine implements those details usually unbeknownst to the client application that generated the high-level document command. Highly reliable database engines typically perform changes to the

1 tables in a transacted manner so that either all the table changes for the particular high-
2 level document command are made, or none of the table changes are made.

3 In addition to implementing the high-level document command in the database (act
4 501), embodiments within the scope of the present invention also include a means or step
5 for ensuring a corresponding notification about the high-level document command is
6 preserved until at least one client application acknowledges the notification. In the
7 embodiment illustrated in Figure 5, this may be accomplished by performing the remaining
8 acts of the method 500.

9 For example, in addition to implementing the high-level document command (act
10 501), the database engine alters the notification table 450 to reflect that the high-level
11 document command has been implemented (act 502). This may be performed after the
12 high-level document command is implemented. However, it is preferred that the
13 notification table be changed atomically with the other tables involved with implementing
14 the high-level document command so that either all the tables are changed together or no
15 changes are made at all. This may be accomplished by having the database engine
16 implement, in a single transaction, the notification table update along with all other table
17 updates needed to implement the high-level document command. If the notification table
18 and the other relevant tables are updated in the same transaction, there is high assurance
19 that the notification table accurately reflects that the high-level document command
20 actually was implemented in the database.

21 The database management system also identifies any client applications that should
22 be notified of the high-level document command (act 503). Typically, this might involve
23 determining which client application are subscribers to the high-level document command
24 although this is not required. A client application could be notified by default without

1 having subscribed or perhaps a third party client application has directed the other client
2 application to receive notifications of the implementation of certain high-level document
3 commands.

4 Once the client applications to be notified are identified (act 503), a notification
5 that the high-level document command has been implemented is then dispatched to the
6 identified client applications (act 504). Messaging may be employed to dispatch the
7 notification as when the database management system and the client application to be
8 notified do not reside on the same machine. Alternatively, function calls through an
9 application program interface or through other interfaces may be employed when the
10 database management system and the client application reside on the same machine.

11 Once notified, the client application may then perform any number and variety of
12 processes in response to receiving the notification. For example, if the original client
13 application has added a new electronic mail message to a public folder, each of the client
14 applications having access to the public folder may be notified. If the client application is
15 already displaying the public folder, the client application would then refresh a display to
16 reflect the addition.

17 Ultimately, the database management system receives acknowledgement back from
18 the client application that received the notification (act 505). The client application may
19 return this acknowledgement as soon as the client application receives notification.
20 However, it may often be beneficial to return acknowledgement well after this time as soon
21 as the client application has completed performing any processes in response to receiving
22 notifications. This ensures that the client application has successfully completed its
23 purpose in being notified in the first place.

24

Once the database management system receives acknowledgement back from the client application (act 505), the database management system alters the notification table 450 to delete reference to the implementation of the high-level document command (act 506). Thus, the notification table reflects the implementation of every high-level document command that has been implemented in the database, but for which the database management system has not yet received acknowledgement back from the client application. In this way, the notification table reflects those notifications that should have been sent to the client application, but for which the notification success is uncertain.

As will now be described, this method permits for notifications of the implementation of high-level document commands to survive even system failures. For example, suppose that, for whatever reason, the high-level document command is implemented in the database, but the appropriate notification is never dispatched. Perhaps the database management system experiences a failure after the high-level document command was implemented, but before the database management system actually sent the notification. In this case, at least the database management system would have set the appropriate entries in the notification table to reflect that the high-level document command was implemented.

Figure 6 illustrates a flowchart of method 600 of how the database management system would recover from such a failure and send the appropriate notification. Upon restarting the database management system (act 601), the database management system checks the notification table to identify any high-level document commands that were implemented, but for which there has not yet been acknowledgement back from the client application (act 602). If there are none (NO in decision block 602), then no notification needs to be resent and the method ends. If there are some (YES in decision block 602),

1 then the subscribers to the notification are identified (act 603), and the appropriate client
2 applications are notified (act 604). Upon receiving acknowledgement back from the client
3 applications (act 605), the appropriate entry in the notification table is cleared (act 606).
4 Thus, even a failure of the database management system will not prevent appropriate
5 notification.

6 Furthermore, if, for some reason, the notification was sent, but the client
7 application never received notification, the client application will still be notified. For
8 example, suppose that the machine hosting the client application was to fail, or perhaps the
9 notification medium (e.g., a network) was to fail resulting in the client application not
10 receiving the notification. In this case, the client application would not return an
11 acknowledgement and the entry would remain in the notification table. The database
12 management system may be configured to continue to send the notification on a periodic
13 basis absent an acknowledgement until acknowledgement is finally returned. Thus,
14 notification may even be made if the client application or network should fail.

15 Thus, the principles of the present invention provide for a reliable way of notifying
16 client applications when a particular high-level document command is implemented in a
17 database. The present invention may be embodied in other specific forms without
18 departing from its spirit or essential characteristics. The described embodiments are to be
19 considered in all respects only as illustrative and not restrictive. The scope of the
20 invention is, therefore, indicated by the appended claims rather than by the foregoing
21 description. All changes which come within the meaning and range of equivalency of the
22 claims are to be embraced within their scope.

23 What is claimed and desired to be secured by United States Letters Patent is:

24